Is Firm Pricing State or Time-Dependent?[†]

Evidence from US Manufacturing

Virgiliu Midrigan[*]

New York University and NBER

March 2008

Abstract

If firm pricing is state, rather than time-dependent, firms are more likely to change prices whenever aggregate and idiosyncratic shocks reinforce each other and trigger desired price changes in the same direction. The distribution of idiosyncratic shocks across adjusting firms therefore varies over time in response to economy-wide disturbances: in times of, say, monetary expansions, the fraction of adjusting firms that have negative idiosyncratic technology shocks should increase. Using measures of technology shocks derived from production function estimates for four-digit US manufacturing industries, I find that sectoral inflation rates are more responsive to negative, as opposed to positive technology disturbances in periods of higher economy-wide inflation, commodity price increases and expansionary monetary policy shocks. I argue, using a quantitative state-dependent sticky price model calibrated to match key features of the US micro-price data, that these results suggest that pricing is state-dependent in US manufacturing.

JEL classifications: E31, E32.
Keywords: state-dependent pricing, time-dependent pricing, technology shocks.

[†]I am indebted to George Alessandria, Bill Dupor, Paul Evans and Mario Miranda for valuable advice and support. I am grateful to Mark Bils, Michael Dotsey, Kevin Huang, Joe Kaboski, Aubhik Khan, and Kei-Mu Yi, as well as several anonymous referees for useful suggestions and comments. All errors are my own.

[*]Department of Economics, New York University, 19 W. 4th Street, 6FL, New York, NY 10012, virgiliu.midrigan@nyu.edu

1. Introduction

Models with nominal rigidities play an important role in recent debates on the role of money in generating business cycle fluctuations, as well as in optimal monetary policy discussions. An important challenge is to build models with solid micro-foundations that are consistent with microeconomic evidence on the price adjustment practices of individual producers, in order to study the aggregate consequences of infrequent price adjustment. The purpose of this paper is to shed additional light on the price-setting practices of firms in the US manufacturing sector. In particular, the question I ask is one that has been at the center of the debate about the potency of nominal rigidities as a monetary transmission mechanism. Is firm pricing state- or time- dependent?

Firms price in a *state-dependent* fashion if the timing of price changes is endogenous, and responds to disturbances to the firm's desired price. State-dependent pricing rules are optimal if the sole friction that prevents price adjustment are physical, menu costs of changing prices and communicating the information about price changes to the consumer. In this environment, an optimal strategy[1] is for firms to follow (S, s) rules: leave their prices unchanged if these are not sufficiently out of line and reprice only in response to a large disturbance (or cumulative history of disturbances). This is the sense in which pricing is state-dependent: the exact date of a price change depends on the current state of the world.

Firms price in a *time-dependent* fashion if the timing of price changes is determined prior to the realization of cost or demand disturbances that affect the firm's desired price. A first generation of time-dependent models, Taylor (1980) and Calvo (1983), postulates that the date at which prices change is outside the control of the firm. In these models firms reprice at exogenously imposed calendar dates even when the potential gains from adjusting in alternative periods are large. The underlying frictions that justify time-dependent rules are, in addition to menu costs, institutional

[1] Seminal contributions include Barro (1972), Sheshinski and Weiss (1977, 1983) in a partial equilibrium context, as well as Caplin and Spulber (1987) and Caplin and Leahy (1991) in general equilibrium.

restrictions, information-gathering or decision-makings costs that render a predetermined schedule of price changes optimal. For example, Zbaracki et. al (2004) report that the pricing season of a large US manufacturing firm occurs once each year, and that new list prices are typically distributed in November of each year. Ball et. al (1988) as well as Bonomo and Carvalho (2004) explicitly model the frictions that render time-dependent pricing rules optimal. In this second generation of time-dependent pricing models the interval between two consecutive price changes is endogenous, and can indeed vary over time in response to changes in the trend growth rate of inflation or the volatility of the environment. Nevertheless, the exact dates at which prices are to be changed are predetermined and independent of contemporaneous disturbances to the firm's price.

The aggregate consequences of the two forms of price rigidities can be very different. Firms' ability to respond to idiosyncratic and aggregate disturbances that are costlessly observed in state-dependent models typically renders monetary policy less potent in these environments. Caplin and Spulber (1987) and Caballero and Engel (1993) show that under special assumptions about the initial distribution of firm prices or shape of the distribution of idiosyncratic disturbances money can indeed be neutral despite nominal rigidities at the firm level. Recent research[2], grounded in explicit household and firm maximization, and using stochastic forcing processes calibrated from the US data, has overturned this neutrality result, but nevertheless reaches the conclusion that state-dependent pricing models generate smaller real effects from monetary shocks.

The key mechanism that leads to smaller real effects from money disturbances in economies with state-dependent pricing is an endogenous shift in the *identity* of adjusting firms that accompanies monetary disturbances. In other words, the endogenous timing of price changes in these models implies that the distribution of idiosyncratic disturbances to adjusting firms' desired prices varies with the aggregate disturbance. With state-dependent pricing firms adjust when they need larger

[2] Dotsey, King and Wolman (1999). Golosov and Lucas (2006).

price changes. This in turn, occurs when the idiosyncratic, say, cost disturbances firms are subject to and the aggregate, say, monetary, disturbances reinforce each other and trigger desired price changes in the same direction. As a result, the mix of adjusters varies with the aggregate shock: in times of positive monetary disturbances adjusters are mostly firms that have received positive idiosyncratic shocks to their desired price; in times of negative monetary disturbances adjusters are mostly firms that have received negative idiosyncratic shocks to their desired price. This change in the mix of adjusting firms imparts a greater degree of flexibility to the aggregate price level as firms that do adjust in a particular period are exactly those firms that desire large price changes of the same sign as the the aggregate disturbance. Golosov and Lucas (2006) call this mechanism the *selection effect* and Caballero and Engel (2007) refer to it as an *extensive margin effect*. This mechanism is very much related to sample selection problems in econometrics: as in those models, the sample of firms that adjusts in any given period in state-dependent economies is non-random (as it would be in, say, a Calvo-type time-dependent environment).

My goal in this paper is to measure the strength of this effect in the US data. A direct measure of the effect aggregate disturbances have on the mix of adjusting firms in a given period requres firm-level data on individual good's prices and idiosyncratic cost or demand disturbances, data that generally unavailable for a large segment of the economy. I rely instead on sectoral price, input and output data available from the NBER Manufacturing Productivity Database. Using this data I first compute measures of technological (cost) disturbances from production function estimates that allow for increasing returns, imperfect competition and variable capacity and labor utilization, using the approach of Basu and Kimball (1997). I then ask whether economy-wide disturbances alter the responsiveness of sectoral inflation rates to these idiosyncratic shocks. The state-dependent model predicts that they should: most adjusting firms are those that have been subject to cost shocks that trigger desired price changes in the same direction as the aggregate disturbance. As a result, the elasticity of sectoral inflation rates to idiosyncratic shocks – my proxy for the fraction of adjusters in a given sector, should increase for those sectors for which the sectoral cost disturbance has the same gradient as the aggregate disturbance.

I indeed find strong support for the hypothesis that aggregate disturbances affect the responsiveness of sectoral inflation rates to sectoral cost shocks in a manner predicted by the state-dependent model. For example sectoral inflation rates are much more responsive to negative, as opposed to positive technology shocks in periods with greater than average aggregate inflation, larger changes in commodity prices and monetary policy shocks. This *selection effect* is both statistically significant and large: it raises the overall response of the inflation rate in the manufacturing sector to a monetary policy shock by up to 50%.

Several earlier papers provide insights into the price-setting practices of individual producers. Blinder et. al. (1998) use survey evidence collected from a survey of CEOs and find that time-dependent rules of price adjustment are twice as common as state-dependent rules. Zbaracki et. al (2004) survey a large manufacturing firm and also find evidence of time-dependent pricing. Cecchetti (1986) and Kashyap (1995) who find that the frequency of price changes increases during periods of higher overall inflation, behavior that is consistent with the implications of state-dependent models, but also of models with endogenous time-dependent pricing. Ball and Mankiw (1994, 1995) illustrate that in menu-cost economies with positive trend inflation, an increase in the volatility of idiosyncratic shocks is inflationary, as most adjusters desire price increases. Similarly, changes in the skewness of the distribution of idiosyncratic shocks can also cause movements in the aggregate price level if pricing is state-dependent. These authors find support for the state-dependent hypothesis as changes in higher-order moments of the distribution of sectoral relative price changes account for an important fraction of changes in aggregate US inflation. Finally, this paper is complementary to that of Midrigan (2006) who studies the strength of the *selection effect* implied by micro-level observations of firm prices in grocery stores. He finds that this effect is muted and money is no longer neutral in a model calibrated to match the fat-tailed distribution of non-zero price changes, as well as the large number of small price changes in the data, but is nevertheless large, as output fluctuations in his setup would be almost twice larger than in the absence of endogenous fluctuations in the identity of adjusting firms.

A final comment on terminology is in order. The terms time- and state-dependent pricing are somewhat obscured by the richness, in recent work, of models with nominal rigidities that employ elements of both of these price setting mechanisms. For example, Ball and Mankiw (1994) formulate a model in which firms price in a time-dependent fashion at predetermined calendar dates, but do have the option of repricing in intermediate periods provided they pay a menu cost. Moreover, discrete-time models of state-dependent pricing, as studied by Golosov-Lucas (2006), Dotsey, King, Wolman (1999) etc. also assume that price adjustment decisions are made at predetermined calendar dates (in periods $t = 1, 2, 3...etc.$) and thus have time-dependent elements. As a result, distinguishing between pure time- and state-dependent models is not the purpose of this paper. Rather, the goal here is to measure the strength of the *selection effect* in the US data. I think of the excersise presented here as providing an additional set of moments useful to enrich our understanding of how firms change prices; not an attempt to reject one model at the expense of another.

The rest of this paper is organized as follows. Section 2 presents a partial equilibrium model with nominal rigidities used to motivate the empirical exercise of this paper. In Section 3 I discuss the data and my measures of technology shocks. Section 4 conducts the empirical analysis. The final section concludes.

2. State- versus Time- Dependent Pricing

In this section I illustrate the selection effect that arises in economies with state-dependent pricing and suggest a set of moments that can be used in order to evaluate the strength of this effect. To do so, I present two widely-used versions of economies with sticky prices in which a) price stickiness arises endogenously, due to menu costs (I refer to this model as one with state-dependent pricing), and b) price stickiness is exogenously imposed, in a Calvo (1983) fashion, and in which the timing and frequency of price changes is exogenous (I refer to this model as one with time-dependent pricing). The two models share many features: I thus present the model economy in a unified fashion and then discuss the differences in the two pricing assumptions.

The model is similar to the partial equilibrium problem studied by Sheshinski and Weiss (1977). A firms' profits depend on its relative price, the ratio of the firm's nominal price to that of the aggregate price level, here assumed exogenous. I assume that the aggregate price level $\bar{p}_t$, evolves according to:

$$\bar{p}_t = \bar{p}_{t-1} e^{g_t}$$

where g_t is the growth rate of the price level and evolves according to

$$g_t = \alpha + \delta g_{t-1} + \eta_t$$

where $\eta_t \sim N(0, \sigma_\eta^2)$. In the discussion that follows I interpret η_t as monetary policy shocks. Let $z_t = \frac{p_t}{\bar{p}_t}$ denote the firm's relative price, where p_t is the firm's nominal price. I assume constant elasticity demand functions: $q_t = z_t^{-\theta}$.

The firm's real profits in period t are: $\Pi(z_t) = z_t^{-\theta}\left(z_t - \frac{c}{a_t}\right)$, where $\frac{c}{a_t}$ is the (real) marginal cost of production, and a_t is the firm's technology[3]. The firm's technology is the product of an idiosyncratic and sectoral component: $a_t = \psi_t \phi_t$, which evolve according to $\log(\psi_t) = \log(\psi_{t-1}) + \varepsilon_t$ and $\log(\phi_t) = \log(\phi_{t-1}) + u_t$, where ε_t is a sectoral and u_t a firm-specific technology shock. The two shocks are drawn from a Gaussian distribution with mean 0 and variance σ_ε^2 and σ_u^2 , respectively[4].

[3] I supress sector and firm subscripts to conserve notation and revert to them below when needed.

[4] The assumption that both marginal cost components follow a unit root is adopted for its computational advantage.

Here, a sector is defined as a group of firms that share the same sectoral technology. I make the distinction between firms and sectors as the empirics is performed on sectoral data. I thus aggregate firm-level decision rules into sectors as described below by computing an average price for firms that share the same sectoral technology ψ_t.

State-Dependent Pricing

In this setup firms face costs of adjusting nominal prices. Specifically, a firm incurs cost ξ every period in which $p_t \neq p_{t-1}$. The firm's problem is to

$$\max_{z_t} E_0 \sum_{t=0}^{\infty} \beta^t \left[\Pi(z_t, a_t) - \xi \mathcal{I} \left(z_t \neq \frac{z_{t-1}}{e^{g_t}} \right) \right], \tag{1}$$

where $\mathcal{I}()$ is an indicator function which takes a value of 1 if the firm adjusts its nominal price.

To write this problem recursively, I bound the state-space so as to ensure that the period reward function is bounded. To this end I define $\hat{z}_t = z_t a_t$. Even though a_t is unbounded, $\hat{z}_t$ is not, as optimality requires z_t be proportional to $\frac{1}{a_t}$. Given this normalization I can write the firm's profits as $\Pi(\hat{z}_t, a_t) = a_t^{\theta-1} \left[\hat{z}_t^{1-\theta} - c\hat{z}_t^{-\theta} \right]$. I can thus rewrite the firm's problem as:

$$V(\hat{z}, g) = \max[V^a(\hat{z}, g), V^n(\hat{z}, g)]$$

where V^a and V^n denote the firm's value of adjusting and not adjusting its nominal price, respec-

Our conclusions are robust to allowing mean-reversion in the technology processes as an earlier version of this paper indicates.

tively that satisfy:

$$V^a(\hat{z}_{-1}, g) = \max_{\hat{z}} \left[\hat{z}^{1-\theta} - c\hat{z}^{-\theta} - \xi + \beta \int\limits_{\varepsilon \times u \times \eta} e^{(\theta-1)(\varepsilon+u)} V\left(\hat{z}'_{-1}, g'\right) dF(\varepsilon, u, \eta) \right], \tag{2}$$

$$V^n(\hat{z}_{-1}, g) = \left[\hat{z}_{-1}^{1-\theta} - c\hat{z}_{-1}^{-\theta} + \beta \int\limits_{\varepsilon \times u \times \eta} e^{(\theta-1)(\varepsilon+u)} V\left(\hat{z}'_{-1}, g'\right) dF(\varepsilon, u, \eta) \right],$$

where $F()$ is the joint cdf of the three shocks. The law of motion for the firm's (normalized) relative price is $\hat{z}'_{-1} = \frac{\hat{z}}{\exp(g)^{\delta}\exp(\eta)} \exp(\varepsilon+u)$ if the firm adjusts its price to $\hat{z}$ and $\hat{z}'_{-1} = \frac{\hat{z}_{-1}}{\exp(g)^{\delta}\exp(\eta)} \exp(\varepsilon+u)$ if it leaves its price unchanged. The term $e^{(\theta-1)(\varepsilon+u)}$ that premultiplies the firm's continuation value is the growth rate of the firm's technology $\left(\frac{a_{t+1}}{a_t}\right)^{\theta-1}$ which enters the profit function in the original problem.

Calvo Time-Dependent Pricing

In this exercise I assume that firms have no control over the timing of their price changes. Rather, the probability that a firm adjusts in a given period is constant, and equal to λ. The functional equations characterizing the firm's problem in this setup are:

$$V(\hat{z}, g) = (1-\lambda)V^a(\hat{z}, g) + \lambda V^n(\hat{z}, g)$$

where V^a and V^n satisfy:

$$V^a(\hat{z}_{-1}, g) = \max_{\hat{z}} \left[\hat{z}^{1-\theta} - c\hat{z}^{-\theta} + \beta \int\limits_{\varepsilon \times u \times \eta} e^{(\theta-1)(\varepsilon+u)} V\left(\hat{z}'_{-1}, g'\right) dF(\varepsilon, u, \eta) \right], \tag{3}$$

$$V^n(s) = \left[\hat{z}_{-1}^{1-\theta} - c\hat{z}_{-1}^{-\theta} + \beta \int\limits_{\varepsilon \times u \times \eta} e^{(\theta-1)(\varepsilon+u)} V\left(\hat{z}'_{-1}, g'\right) dF(\varepsilon, u, \eta) \right],$$

To solve these problems, I employ collocation, a functional approximation technique. The idea behind this method is to approximate the two value functions with a linear combination of orthogonal polynomials and solve for the unknown coefficients by requiring that the two equations are satisfied exactly at a number of nodes along the state-space. A technical appendix discusses the solution method and its accuracy in more detail.

Calibration

I assign the model parameter values to ensure that the predictions of the state-dependent model match certain features of the US economy. The length of the period is set to one quarter, but the model will be evaluated against annual data from the NBER Productivity database. I therefore choose the parameters that characterize the process for the growth rate of the aggregate price level to match the annual mean, serial correlation and volatility of inflation in the manufacturing sector. This gives $\alpha = 0.00098$, $\delta = 0.89, \sigma_\eta^2 = 2.5 \times 10^{-5}$. The elasticity of demand, θ, is chosen so that the steady-state markup is equal to 25%[5]. This leaves three additional parameters that must be calibrated in the state-dependent economy: ξ, the menu costs, as well as the volatility of sectoral σ_ε^2 and firm-specific σ_u^2 technology disturbances. The three targets that pin down these parameters are a) an average duration of contract lengths of 15.3 months (frequency of quarterly price changes of 0.196[6]), which corresponds to an estimate of Leith and Malley (2007) regarding the frequency with which firms in the NBER productivity database change prices, b) an average size of non-zero price changes of 9%, consistent with findings from Nakamura and Steinsson (2007) and Bils and Klenow (2005), and c) a standard deviation of annual sectoral inflation rates of 4.9% in the sectoral price data available from the NBER Productivity Database. Finally, notice that in the absence of menu

[5] This number is in line with elasticity of substitution estimates in earlier work, which range from $\theta = 3$ to $\theta = 6$. See Obstfeld and Rogoff (2000) for a brief survey.

[6] An earlier version of this paper targets a frequency of price changes of 0.33 per quarter, consistent with recent evidence from micro producer prices by Nakamura and Steinsson (2007). Less frequent price adjustment makes it easier to identify the selection effect at the annual frequency and brings the model's quantitative predictions more in line with the data.

costs ($\xi = 0$) in which case the firm's value V is flat in its two arguments, the firm's discount factor is $\beta E e^{(\theta-1)(\varepsilon+u)} = \beta e^{\frac{(\theta-1)}{2}(\sigma_\varepsilon^2+\sigma_u^2)}$. I choose β to ensure that the firm's discount factor corresponds to an annual real interest rate of 4%. I thus set $\beta = \frac{0.96^{\frac{1}{4}}}{e^{\frac{(\theta-1)}{2}(\sigma_\varepsilon^2+\sigma_u^2)}}$. The table below summarizes the parameter values I use. The menu cost, ξ, is equal to 1.44% of the firm's steady-state revenues. Firm-specific shocks are 3 times more volatile than sectoral shocks.

β	α	δ	σ_η^2	σ_ε^2	σ_η^2	ξ	θ
.97	2.8×10^{-3}	0.76	2.32×10^{-5}	1.95×10^{-4}	6.47×10^{-4}	0.0144	5

As for the Calvo model, I choose λ to match a frequency of price changes of 0.196, as in the state-dependent model, and assign the same volatility of technology shocks.

Optimal Pricing Rules

As is typical in economies with menu costs, firms follow generalized (S, s) rules and reprice whenever shocks, whether aggregate, sectoral, or idiosyncratic, force $\hat{z}_{-1}$ to drift away from its optimum (which turns out to be close to $\frac{\theta}{\theta-1}c$, the frictionless optimum). Figure 1 illustrates the firms' value of changing and that of not changing its price as a function of $\log\left(\hat{z}_{-1}/(\frac{\theta}{\theta-1}c)\right)$: the deviation of the firm's price from its frictionless optimum. If a firm adjusts its price, is value is independent of z_{-1}, by inspection of the firm's problem in (2). In contrast, if the firm does not adjust its price, it sells at z_{-1} and the further away z_{-1} is from its optimum, the lower the firm's value. The intersection of these two value functions determines the firm's inaction and adjustment regions; whenever $\log\left(\hat{z}_{-1}/(\frac{\theta}{\theta-1}c)\right)$ is sufficiently away from zero the firm finds it worthwile to pay the menu cost and adjust. In contrast, the Calvo firms' adjustment decisions is exogenous: firms adjust with a constant hazard $1 - \lambda$.

Consider next the firms' pricing rules. As the recursive representation of the problem indicates, in the presence of a unit root in the process for technology the firm's price, conditional on adjustment, depends solely on the growth rate of the aggregate price level. Given that this growth rate, g, is persistent, it helps forecast future changes in the growth rate of the price level and therefore affect the adjusting firm's optimal price. As Figure 2 indicates Calvo firms respond more aggressively to an increase in the growth rate of the price level than state-dependent firms do. These differences in price functions arise because of the type of nominal frictions Calvo and menu cost firms are subject to. If a Calvo firm finds itself with a suboptimal price in a given period in the future, it pays dearly: given that it will not re-adjust its price for an average of 3 quarters, it will incur losses from the suboptimal price for a number of periods to come. In contrast, a state-dependent firm can always choose to pay the menu cost and reprice: its losses from having a suboptimal price in future periods are smaller than those of a time-dependent firm. This in turn implies that a Calvo firm has a stronger incentive to offset future expected changes in its marginal cost every time it adjusts than a state-dependent firm does. A similar argument explains why Calvo firms choose higher prices on average than state-dependent firms do: they have a stronger incentive to respond to the trend growth in the aggregate price level, as captured by α.

How do firms respond to technology disturbances? The fact that adjusting firms choose to return their normalized relative price to a target level $\hat{z}^*$, which depends on g only implies that the firm's nominal price p responds one-for-one to a technology disturbance both in the Calvo and menu-cost world. Thus an increase in a, either because of an increase in the sectoral technology ψ or the idiosyncratic technology ϕ leads firms to lower their nominal prices one-for-one[7].

I next aggregate firm decision rules into sectors in order to derive several moments of the sectoral price data that can be used empirically. Figure 3 plots the fraction of adjusters $\mathcal{F}r(\varepsilon, g)$ in a sector that starts at the ergodic distribution of prices and is then subject to a sectoral technology shocks ε, and a growth rate of the price level equal to g, in the state-dependent model. In addition,

[7]Lowering the serial correlation of technology reduces the responsiveness of firm prices to technology shocks as these are expected to die out in future periods. Calvo firms respond less than menu-cost firms, but the elasticities of prices to technology shocks are approximately constant in both economies.

firms are subject to idiosyncratic technology shocks u. This fraction is, by assumption, constant in a time-dependent economy like Calvo.

Consider first sectors subject to negative technology disturbances ($\varepsilon = -.06$ and $\varepsilon = -.03$). Firms in these sectors desire, on average, to increase their prices in order to respond to the higher marginal costs of production. This incentive to change prices is reinforced if the economy-wide nominal shock is also positive. For this reason, the fraction of firms that adjusts in sectors with negative technology shocks increases in g, the growth rate of the general price level.

In contrast, firms in sectors with positive technology shocks see their marginal cost falling and desire price decreases. Their desire to decrease real prices is automatically satisfied if the aggregate price level increases, thereby eroding the firm's real price. The fraction of firms that adjusts in sectors with positive technology disturbances therefore decreases as the growth rate of the economy-wide price level rises.

The results presented in Figure 3 are not useful for empirical purposes, as I do not directly observe the fraction of firms that adjust in a given sector. This fraction is, however, well proxied by the elasticity of a sector's inflation rate to its technology disturbance. Given that adjusters respond one-for-one to a sectoral technology shock, the elasticity of sectoral inflation rates is correlated with the fraction of firms adjusting prices. To see this, suppose that adjusting firm j in sector i chooses an inflation rate in period t that is equal to:

$$\pi^*_{ijt} = \varepsilon_{it} + g_t + \tilde{u}_{ijt},$$

which, in light of our discussion above, is a good approximation to a menu-cost firms' optimal price function. Here $\tilde{u}_{ijt}$ captures the idiosyncratic shock to the firm's desired price and includes the contemporaneous cost shock u_{ijt} as well as the cumulative history of idiosyncratic and aggregate disturbances since the firm has previously adjusted. Letting $\Theta(\varepsilon_{it}, g_t)$ denote the firm's adjustment region in the u_{ijt} space (the set of u_{ijt} for which a firm adjusts), the sectoral inflation rate, defined as $\pi_{it} = \int \pi_{ijt} dj$ is then equal to

$$\pi_{it} = \mathcal{F}r(\varepsilon_{it}, g_t)\,(\varepsilon_{it} + g_t) + \int\limits_{\tilde{u}_{ijt} \in \Theta(\varepsilon_{it}, g_t)} \tilde{u}_{ijt} dj$$

The second term in this expression captures the selection effect at the sectoral level: firms with $\tilde{u}_{ijt}$ alligned with $\varepsilon_{it} + g_t$ are more likely to adjust prices. Thus, although a regression of π_{it} on $\varepsilon_{it} + g_t$ provides an upward biased estimate of $\mathcal{F}r(\varepsilon_{it}, g_t)$ because of the selection bias, the regression coefficient is nevertheless correlated with the fraction of adjusting firms.

Given this discussion, one way to test the state-dependent model is to estimate, for each period t, the following cross-sectional regressions of sectoral inflation rates on sectoral technology shocks[8] in which the coefficient on negative, (γ_t^N) and positive, (γ_t^P) shocks, are allowed to differ:

$$\pi_{it} = \xi_t + \gamma_t^N \varepsilon_{it} \mathcal{I}_{\varepsilon_{it} < 0} + \gamma_t^p \varepsilon_{it} \mathcal{I}_{\varepsilon_{it} > 0} + u_{it}$$

If pricing is indeed state-dependent, γ_t^N should increase in absolute value (become more negative) in periods in which the economy is experiencing an aggregate shock that raises all firms' prices as sectors in which are firms are concomitantly hit by a negative sectoral technology shock have more firms adjusting. Similarly, γ_t^P should fall in absolute value (become less positive in periods in which the economy is hit by a positive aggregate shock). To illustrate this, I aggregate individual firm decision rules and simulate the state- and time-dependent economies above and compute these elasticities for which period. Figure 4 plots γ_t^N and γ_t^p against the aggregate disturbance, g_t across periods,

[8] I could allow for non-linearities in the elasticity of sectoral inflation rates to aggregate shocks as well, but the fact that technology shocks, in both the model and the data, are much more volatile than aggregate disturbances makes them a more suitable proxy for the fraction and identity of firms that adjust in a given sector.

for the state- and time- dependent models. As anticipated, the absolute value of γ_t^N increases with the aggregate shock: when g_t is close to 0, the elasticity is close to 0.75 and as g_t increases to 0.03 this elasticity rises to 1 in absolute value. Similarly, the absolute value of γ_t^P falls from 0.9 to 0.7 as g_t increases from 0 to 0.03. Notice also that γ_t^P is somewhat flatter in g_t than γ_t^N is and that this slope flattens as g_t increases. This is an outcome of the trend growth rate in the aggregate price level which implies that in simulations g_t is mostly positive. As Figure 3 indicates, the fraction of adjusting firms for sectors with positive shocks is flatter than for sectors with negative shocks in the region in which g is positive as in this region the aggregate shock is cancelled by the sectoral shock. This is a typical feature of state-dependent models: the hazard of price changes is flatter for smaller deviation of the desired price from its target than for larger deviations[9]. Also note that, if the aggregate shock increases even further, firms in all sectors of the economy, irrespective of their sectoral disturbances, would find optimal to increase their prices. Thus, given that the fractions in Figure 3 or elasticities in Figure 4 are drawn for small values of the aggregate disturbance, as in the US data, all statements above hold locally, rather than globally.

In practice, the two-stage procedure outlined above of a) estimating elasticities for each period using cross-sectional regressions and b) relating these elasticities to measures of aggregate disturbances is inneficient. One can instead parameterize γ_t^N and γ_t^P directly as functions of the aggregate shock:

$$\begin{aligned} \gamma_t^N &= \beta_0 + \beta_1 g_t, \\ \gamma_t^P &= \alpha_0 + \alpha_1 g_t, \end{aligned}$$

and infer the size of the coefficients β_1 and α_1 from panel regressions that pool observations across

[9] See, e.g., Caballero and Engel (2007).

sectors and time-periods together:

$$\pi_{it} = \xi_i + \gamma_t^P \mathcal{I}_{(\varepsilon_{it}>0)}\varepsilon_{it} + \gamma_t^N \mathcal{I}_{(\varepsilon_{it}<0)}\varepsilon_{it} + \rho g_t + u_{it}$$

where ξ_i are sector-specific fixed effects that account for differences in the trend growth rate of prices across sectors.

As earlier, if pricing is state-dependent, higher aggregate inflation, g_t, should increase the fraction of adjusting firms in sectors in which the technology shocks are negative and decrease the fraction of adjusers in sectors in which technology shocks are positive. As a result, higher aggregate inflation should increase the (absolute value of, i.e., make it more negative) elasticity of sectoral inflation rates to sectoral technology shocks, in sectors in which technology shocks are negative and decrease (in absolute value, i.e., make it less negative) in sectors with positive technology shocks. The state-dependent model thus suggests that $\beta_1 < 0$ and $\alpha_1 > 0$. Substituting out the definitions of γ_t^N and γ_t^P, the regression I propose to estimate is:

$$\pi_{it} = \xi_i + \alpha_0 \mathcal{I}_{(\varepsilon_{it}>0)}\varepsilon_{it} + \alpha_1 \left(g_t \times \mathcal{I}_{(\varepsilon_{it}>0)}\right)\varepsilon_{it} + \beta_0 \mathcal{I}_{(\varepsilon_{it}<0)}\varepsilon_{it} + \beta_1 \left(g_t \times \mathcal{I}_{(\varepsilon_{it}<0)}\right)\varepsilon_{it} + \rho g_t + \varsigma_{it}. \quad (4)$$

In Table 1 I report the coefficient estimates in this regression computed using model-simulated data. I construct sectoral inflation rates in a manner that attempts to mimic the nature of the empirical data I study in the next section. In particular, I simulate firm-decision rules for 446 sectors, as in the NBER Productivity Database, for 36×4 quarters, as 36 years of data are available for empirical analysis. Each sector is made up of 125 firms: this number is the weighted (according to each sector's sales share) average of the number of firms in each SIC-4 digit manufacturing sector as measured by the inverse of the Herfhindhal-Hirschmann concentration ratio for 1992 reported by the US Census Bureau[10]. The period in the model is one quarter, however, the NBER Productivity Database reports annual sectoral observations. To allow comparison between the model and the data I construct annual sectoral inflation rates, π_{it}, by computing a Paasche index using price and quantity data for individual firms. To compute sectoral productivity shocks, I divide total output produced in a given year by all firms in an industry by their total labor input. I use changes in log of this measure of labor productivity as a measure of industry-specific shocks, ε_{it}. Price stickiness at the firm level makes this is an imperfect measure of the ε_{it} given that output is demand determined, but by mimicking the empirical exercise of the next section I can quantitatively compare coefficient estimates in the model and in the data[11]. Finally, g_t at the annual level is constructed by summing consecutive 4-quarter non-overlapping sets of the quarterly growth rate of the model's exogenous driving process, g_t, and filtering out low-frequency variations using an HP(10) filter in the spirit of the business cycle literature. The consequences of filtering (and estimates without filtering) are discussed in the data section below, but shortly, filtering increases the strength of the selection effect, both in the model and in the data, presumably because it allows us to isolate unexpected shocks to inflation that are not yet incorporated in the firms' prices. I employ 500 rounds of model simulations and report means and standard deviations of coefficient estimates across the different simulations in parantheses. I repeat this exercise for both the state- and time-dependent models.

Time aggregation to annual data clearly washes out some of the non-linearities reported in Figures 3 and 4 at the quarterly frequencies. To the extent to which most firms are able to respond to aggregate and sectoral shocks from one year to another, the importance of price stickiness, whether time- or state-dependent is reduced. As a result, as the first column (State-Dependent) of

[10]The number of firms in each sector is much larger, because of size heterogenity across firms in a given sector. In the absence of size dispersion, the inverse of the HH concentration ratio is equal to the number of firms in that sector, and I use this measure to control for size dispersion.

[11]An earlier version of this paper uses the actual ε_{it} (summed across 4 quarters) measures to estimate the coefficients in the above equation. Although qualitatively similar, these coefficients were much smaller in absolute value, suggesting that the absolute value of the coefficient estimates are sensitive to the details with which data is time-aggregated.

Table 1 illustrates, the coefficient on positive technology shocks, α_1 is insignificantly different from 0 (although negative, the mean across 500 simulations is twice less than the standard deviation) at the annual frequency. In contrast, the avearge coefficient on negative shocks is large in absolute value and more than twice larger its standard deviation. The intuition for why the coefficient on positive shocks is virtually zero is the same as in our discussion of Figures 3 and 4 above and in Ball and Mankiw (1994): firms in sectors with positive shocks to their technology are less willing to change prices given that their incentive to lower the price is offset by trend aggregate inflation, and are thus in a flatter region of their adjustment hazard, making their elasticities to technology shocks less responsive to aggregate inflation.

In contrast to the State-Dependent model, all coefficients on the interaction terms are insignificantly different from zero in my simulations of the Time-Dependent model, consistent with the evidence in Figure 4. Together, the two columns of Table I suggest that one can measure the strength of the selection effect in the data using estimates of equation 4. Finally, notice that simulations of the state-dependent model consistently produce elasticities of inflation rates to aggregate (g_t) and sectoral (ε_{it}) shocks that are greater than those in the time-dependent models. This is because of the selection effect at the individual firm level and at the sectoral level (that is not completely soaked up by our non-linear terms).

3. Data

I test the predictions of the state-dependent pricing model using annual data from 1958 to 1996 for 446 4-digit SIC industries from the NBER Manufacturing Productivity Database[12]. The data is derived from various government sources, notably the Census Bureau's Annual Survey of Manufacturing, and contains information on total shipments, materials expenditure, investment,

[12]The data is available at http://www.nber.org/nberces/nbprod96.htm and is discussed extensively in Bartelsman and Gray (1996). The industries are those defined in the 1972 Standard Industrial Classification. We drop two industries that have missing observations for several years.

capital stock, number of production and non-production workers, payroll, production worker hours and wages, as well as price deflators for shipments, materials etc. for each industry. Material expenditures include expenditure on energy, and the deflator for materials accounts for movements in the price of energy. Bils and Chang (1999) is a recent example that uses this dataset in order to ask how industry prices respond to variations in costs and production, although, given my focus on asymmetries in response to purely technological shocks, my approach differs from theirs along several dimensions. I use this data in order to conduct my empirical exercises as discussed below.

A. Measuring technology shocks

My measures of technology shocks are Solow (1957) residuals estimated using the methodology developed by Hall (1990) and Basu and Kimball (1997) in order to account for the possibility of increasing returns, imperfect competition and variable input utilization, respectively.

I assume a differentiable production function in which firms produce output Y, using capital services $\tilde{K}$, labor services $\tilde{L}$, intermediate inputs of materials and energy M according to:

$$Y = F(\tilde{K}, \tilde{L}, M, A)$$

Capital services depend on the stock of capital K, but also capital utilization $Z : \tilde{K} = ZK$, while labor services depend on the number of workers N, hours worked per employee H and each worker's effort level $E : \tilde{L} = ENH$. Taking logarithms of this production function, totally differentiating, and invoking cost minimization, one obtains:

$$dy = \mu \left[s_k dk + s_L(dn + dh) + s_m dm \right] + \mu \left[s_k dz + s_L de \right] + da$$

where lower case letters denote logs, s_j is the share of factor j in total revenue and μ is the markup . The difficulty in estimating this equation directly is that effort and capital utilization are not observed. I follow Basu and Kimball (1997) and proxy the unobserved input utilization with hours per worker dh[13]. The justification for this approach is that firms operate along all margins simultaneously, and given convex costs of changing hours worked, effort and capital utilization, will choose to change them simultaneously in response to a shock. Changes in hours worked are therefore correlated with unobserved capital utilization and effort. More formally, Basu and Fernald (2000) solve a dynamic cost minimization problem of a firm subject to costs of changing employment levels, hours worked and capital utilization, and show that as long as capital's depreciation rate does not depend on its utilization level and the production function is Cobb-Douglas, a log-linear approximation to the firm's optimality conditions implies that dz and de depend on dh only[14]. I therefore estimate

$$\Delta y_{it} = c_i + \mu \Delta x_{it} + \gamma \Delta h_{it} + \tilde{\varepsilon}_{it} \tag{5}$$

where Δy_{it} is the change in the log output of industry i, Δx_{it} is the share-weighted sum of the growth rate of real inputs (labor, capital, materials and energy). I calculate total output as shipments plus change in end-of-period inventories and deflate it using the price deflator for shipments. The Productivity Database distinguishes between production and non-production workers in reporting industry employment, and only reports hours data for production workers. I use the two as separate inputs in the production function and assume that hours per worker are time-invariant for non-production workers. My results are robust to an alternative measure of inputs that includes only production workers. My proxy for variable input utilization, Δh_i, is the log-difference in hours per

[13]Conley and Dupor (1999) use an alternative proxy for capital utilization, one based on electricity consumption. Electricity data is not available however at the 4-digit level of disaggregation.

[14]Allowing for depreciation rates to increase with capital utilization, as in Basu and Kimball (1997) complicates the problem as utilization will depend on material inputs, capital stock, investment and the relative price of materials and investment: $dz = A(d(p_m - p_I) + dm - dk) + B(di - dk)$. where $p_m - p_I$ is the relative price of materials and investment, i and k are investment and the stock of capital, respectively, A and B are constants. We have used this alternative proxy for capital utilization and found results to be very similar to those reported in text.

worker reported for production workers.

I calculate the share of each factor of production as the time-series average of total payments to each factor divided by total revenues in each industry. One could in principle depart from this Cobb-Douglas assumption of constant shares and allow shares to vary over time, but, as Basu and Fernald (2000) argue, this approach increases the likelihood of misspecification because observed factor prices are not allocative period-by-period in a world with implicit contracts or quasi-fixity[15]. To calculate payments to capital, I first calculate the user cost of capital, R, according to[16]:

$$R = (r + \delta)\frac{1 - ITC - \tau d}{1 - \tau}$$

where r is the required rate of return on capital (I follow Hall (1990) and assume it equal to the S&P 500 dividend yield), δ is the depreciation rate, ITC is the investment tax credit, d is the present-value of depreciation allowances and τ is the corporate income tax rate. Jorgenson and Yun (1991) provide data on ITC, d and δ for 53 types of capital goods, while the tax data is provided by the Bureau of Economic Analysis at the 2-digit level of disaggregation. I calculate the user cost of capital for each asset and a weighted average over the different types of assets for each SIC 2 industry in the dataset, with the weights reflecting the relative importance of each type of asset in each industry. I judge the relative importance of the different types of assets in each industry by using Bureau of Economic Analysis data on the 1982 Distribution of New Structures and Equipment to using industries. The required payment to capital is finally calculated as RP_kK where P_kK is the current-dollar value of the industry's stock of capital[17]. Given that the Database only reports wage and salary costs of labor, I follow Bils and Chang (1999) and magnify both production and

[15] Our results are robust however to an alternative specification in which factor shares vary over time and are equal to average shares in adjacent periods.

[16] See Hall and Jorgenson (1967).

[17] We also follow Bils and Chang (1999) and, given the low level of profits in manufacturing, calculate capital's share residually, assuming that the shares of all inputs sum to one. Results are robust to this alternative assumption.

non-production labor costs to account for employer pension payments and compensation benefits. This data is again based on information available in the underlying NIPA tables at the 2-digit level of disaggregation. In addition, I magnify total labor costs (for both production and non-production workers) by 9% to account for the database's exclusion of payments to auxiliary and support personnel. Bartelsman and Gray (1996) report that these costs account for 7.9% and 10.7% of total payroll in manufacturing in 1972 and 1986 respectively.

OLS estimates of (5) are likely to be biased because of the correlation between technology shocks and input choices. I therefore instrument the right-hand side variables using current and one period lags of deflated oil price changes, changes in government spending, changes in the US effective nominal exchange rate and monetary policy shocks estimated using a 7-variable VAR according to the Christiano, Eichenbaum and Evans (1999) block-recursive identification procedure[18]. My instruments are similar to those used by Basu and Kimball (1997), to which I add a measure of changes in nominal exchange rates of US against its trading partners. Given the exchange rate disconnect puzzle documented in open-economy macroeconomics[19], it is unlikely that sectoral technology shocks are correlated with this variable[20].

The relatively short span of time-series observations renders industry by industry estimates of the coefficients in (5) rather imprecise. I therefore pool 2-digit industries together and estimate (5) using a panel (fixed-effects) 2SLS estimator for each SIC 2 industry[21].

Although my interest is not in estimates of equation 5 per se, I briefly compare my estimates to those in earlier work. The time-series standard deviation (average across all sectors) of the purified Solow residuals is 0.063, whereas that of changes in the TFP (the difference between the

[18] We measure the stance of monetary policy by the size of non-borrowed reserves and assume that the Fed's information set includes current and four lagged values of real GDP, CPI, an index of commodity prices, as well as four lags of the Federal funds rate, total reserves, non-borrowed reserves and the M1 money stock. Monetary policy shocks are estimated using quarterly data. Our measure of annual shocks is the sum of four quarterly shocks.

[19] e.g, Obstfeld and Rogoff (1996)

[20] Our results are robust to excluding nominal exchange rate variables as an instrument for input growth.

[21] Our results are robust however to estimating technology shocks using separate industry-by-industry regressions.

growth rate of real output and hare-weighted sum of the growth rates of inputs) is 0.076. The two series are strongly correlated (0.63). In contrast, Basu and Fernald (2000) estimate that technology shocks in the entire manufacturing sector are almost twice less volatile: the standard deviation of the Solow residual is 0.035 and that of the purified series is 0.028 according to their estimates. My estimates of returns to scale are also similar to those of Basu, Fernald and Kimball (2004) who use the Jorgenson dataset of 29 industries (including 21 industries at (roughly) the SIC 2 level) from 1949 to 1996. Their estimation strategy differs slightly from mine as they restrict the coefficient on the proxy for input utilization to be constant across industries, but, despite the differences in the level of aggregation underlying the two sets of estimates, my results and theirs are not too dissimilar. For durable goods, the median returns to scale estimate is 1.11, compared to 1.07 in their work, with a correlation of 0.71 across coefficient estimates in the different industries. For non-durables, the correlation is 0.77 if one excludes two industries (food and leather) for which these parameters are imprecisely estimated, while the degree of returns to scale is higher in my work (1.07), than theirs (0.89)[22].

4. Empirics

Before I discuss formal estimates from panel regressions, I first use my estimates of technology shocks $\varepsilon_{it} = \tilde{\varepsilon}_{it} + c_i$ constructed above[23], to estimate the following cross-sectional regressions:

$$\pi_{it} = \xi_t + \gamma_t^N \varepsilon_{it} \mathcal{I}_{\varepsilon_{it}<0} + \gamma_t^p \varepsilon_{it} \mathcal{I}_{\varepsilon_{it}>0} + u_{it} \tag{6}$$

[22]My estimates are however much closer to those of Burnside, Eichenbaum and Rebelo (1995) who find returns to scale in non-durable manufacturing equal to 1.13.

[23]Throughout this paper, my definition of a sector's technology shock is the sum of the "purified Solow" residuals in equation (10) plus the fixed-effect term that captures the long-run rate of growth of an industry's technology. Results are little affected if I use instead the residuals themselves as a measure of technology shocks.

for each time-period using ordinary least squares, where π_{it} are sectoral inflation rates. Recall (Figure 4) that the state-dependent model predicts that γ_t^N should increase in absolute value and γ_t^P should fall in absolute value in periods of higher aggregate disturbances. In Figure 5 I plot the (absolute value of) two elasticities against three measures of aggregate shocks that increase firms' desired prices: HP-filtered inflation in the manufacturing sector, commodity price changes, and a (2-year lagged) measure of monetary policy disturbances described above. The model is silent as to what measures of aggregate shocks should be included: any disturbance that affects all firms' desired nominal prices should raise the fraction of adjusters in sectors with negative shocks. Commodity price changes and monetary disturbances are natural proxies for g_t in the model. So is manufacturing inflation, as in the model firm's prices on average respond strongly to the nominal disturbance g_t, for arguments discussed in Caplin-Spulber (1987), Golosov-Lucas (2006) and Midrigan (2006).

The line through these scatter plots is from a fitted OLS regression of $\gamma_t^N(\gamma_t^P)$ against the variable on the x-axis. For all 3 measures of shocks the elasticity on negative sectoral shocks increases with the aggregate disturbance. Similarly, the elasticity on positive sectoral shocks decreases with commodity price inflation and monetary policy shocks, although it does not vary with HP-filtered inflation. This last fact should not be of much concern as simulations of the model reported in Table 1 do not suggest a strong relationship between the elasticity on positive shocks and the aggregate disturbance at annual frequencies.

I next proceed to formally measuring the size of the selection effect in the US data employing a panel specification in which elasticities to positive and negative shocks are directly parameterized as functions of aggregate disturbances. In particular, I employ the same regression specifications as in (4), which I repeat here for convenience:

$$\pi_{it} = \xi_i + \alpha_0 \mathcal{I}_{(\varepsilon_{it}>0)}\varepsilon_{it} + \alpha_1 \left(g_t \times \mathcal{I}_{(\varepsilon_{it}>0)}\right)\varepsilon_{it} + \beta_0 \mathcal{I}_{(\varepsilon_{it}<0)}\varepsilon_{it} + \beta_1 \left(g_t \times \mathcal{I}_{(\varepsilon_{it}<0)}\right)\varepsilon_{it} + \rho g_t + \varsigma_{it}.$$

As in Figure 5 I use three alternative measures of aggregate disturbances, g_t : manufacturing inflation (HP-filtered), commodity price inflation, as well as a measure of monetary policy shocks due to Christiano, Eichenbaum and Evans (1999) which uses a recursive identification assumption and non-borrowed reserves as the postulated instrument[24]. I use these alternative measures of shocks, instead of focusing solely on CPI inflation in order to ascertain the robustness of my results but also in order to establish causality as exogenous variations in γ_t^N and γ_t^P can themselves trigger variation in aggregate inflation. In particular, the two elasticities can fluctuate, and thereby affect inflation in the presence of exogenous changes in higher-order moments of the distribution of idiosyncratic cost shocks, as in Ball and Mankiw (1995).

I present the results in Table 2[25]. For comparison, the first column of the table reports the size of these coefficients predicted by the state-dependent model, together with the standard deviation, across different simulations of the model, of these coefficients in parantheses. The next 4 columns, labeled I-IV, correspond to the different measures of g_t in the regressions. I report, in parantheses, standard errors for these coefficients. These standard errors are corrected for the bias that arises because of our use of a two-stage procedure that imparts uncertainty to our estimates of the technology shocks, ε_{it}[26], as well as for heteroskedasticity and serial correlation across industries of arbitrary form, by employing a Arellano (1987)-type correction. I describe the approach used to correct for the two-stage bias in the appendix.

Notice first in columns I-III of Table 5 that sectoral inflation rates increase with all measures of aggregate shocks and decrease with sector-specific technology shocks. The size of the coefficient on

[24] An earlier draft of this paper has used two additional measures of monetary disturbances due to Romer and Romer. These alternative measures produce similar results. To construct a measure of annual money shocks, I first estimate them at the quarterly frequency, and then sum up these shocks for the 4 quarters during a year.

[25] The regression with monetary policy shocks includes current and 1-, 2-, and 3-year lagged monetary shocks. I only report the coefficients in the period in which money has the largest effect on sectoral inflation to conserve space. In an earlier version of the paper I have shown however that results are consistent with the state-dependent model for lags 1-3, and go in the wrong direction in the period of the shock, this latter result reminiscent of the well-documented "price puzzle" in earlier work.

[26] In the appendix I describe the exact correction employed.

the aggregate shock, g_t, varies substantially from one specification to another, which is not surprising, given that these alternative measures are characterized by different degrees of persistence and the fact that firms in a sticky price environment respond more aggressively to more persistent disturbances which are expected to last longer. The size of the coefficient on technology shocks is similar across the last three columns, and always of a negative sign. The absolute value of these coefficients is, however, smaller in the data, which is perhaps evidence of strategic complementarities that prevent firms from fully responding to sectoral shocks[27]. Alternatively, there may be more mean reversion in the process for sectoral technology shocks than imposed by our unit root assumption in the model. Finally, notice that the coefficient that captures most strongly the *selection effect*, β_1, is, of the sign and magnitude predicted by my simulations of the state-dependent model. Firms in sectors with negative technology shocks raise prices faster in periods of positive aggregate disturbances, suggesting state-dependence in their pricing decisions. As seen in Table 1, the model suggests that the non-linear response to positive shocks, as measured by α_1 should wash out at the annual frequency. The data produces mixed implications regarding the size of this effect. Using manufacturing inflation as a measure of disturbances I estimate a coefficient $\hat{\alpha}_1$ that is indistinguishable from 0 statistically. Using commodity price changes and monetary policy shocks I obtain positive coefficients, in line with the predictions of the model at quarterly frequencies (Figure 4). What is more important for the aggregate consequences of these non-linearities is that for all three measures of shocks, the elasticity on negative shocks rises (in absolute value) relative to that on positive shocks, thereby increasing the flexibility of the aggregate price level in excess of that in time-dependent models that assume a hazard of price adjustment independent of the size of the shock.

In Table 3 I report a robustness check in which I use actual US manufacturing inflation as a

[27] One (crude) way of capturing strategic complementarities in my partial equilibrium setup is to allow for convex costs of price adjustment that penalize firms for large price changes. Adding these costs to the model lowers the estimated coefficients α_1 and β_1 in model simulations. Results from these simulations are available from the author upon request.

measure of aggregate disturbances. I report results for filteted and unfiltered measures of inflation, both in the model and in the data. First, notice that the model simulations show a stronger selection effect (as measured by the absolute value of β_1 (5.39 vs. 3.63) or the difference between β_1 and α_1 (4.59 vs. 2.72) when HP-filtered inflation is used as a measure of aggregate disturbances. An argument in Ahlin and Shintani (2007) suggests that this is the case because in environments with persistent inflation transition from a period of high inflation (in which firms charge higher prices in expectation of high inflation in future periods) to low inflation (in which firms charge lower prices as the incentive to front-load future desired price increases is reduced) may lower firms' desired prices. HP-filtered measures of inflation may thus account for this effect by conditioning the elasticities on deviations of inflation from its trend and thus controlling for the effect of past inflation on firms' desired prices.

As in the model, use of actual inflation provides stronger evidence of a selection effect in the data. The absolute value of β_1 drops from 6.56 to 4.49, while the difference between the two coefficients drops from 6.29 to 2.52. The value of α_1 (-1.97) is negative and significantly different from zero, which is somewhat inconsistent with the model's intuition, but recall that the model has no tight implications regarding the sign of α_1 at the annual frequency. The large standard deviation of estimates of α_1 in simulations puts the coefficient of -1.97 in the data well in the range of the estimates in the model.

A final robustness check I perform in the last column of Table 3 is to instrument inflation to correct for potential endogeneity of manufacturing inflation, which is itself endogenous to the elasticties on sectoral shocks. I instrument inflation with oil price changes, monetary policy shocks and changes in the nominal exchange rates. Results are very similar to those reported in the adjacent column and if anything, suggest an even stronger selection effect. In particular, the coefficient on positive shocks is now virtually 0.

A. Interpreting the results

I have established above the statistical significance of state-dependent pricing terms in explaining fluctuations in sectoral inflation rates. I next ask whether their effect is quantitatively large. I use my estimates of equation (4) and calculate, in Table 4, the effect of a one standard deviation increase in the different measures of aggregate shocks on the elasticity of sectoral inflation to negative/positive technology shocks. I first calculate what the elasticities γ_t^N and γ_t^P would be in the absence of economy-wide disturbances, when the aggregate variables are at their time-series means: these are the estimates of $\alpha_0 + \alpha_1 \times mean(g)$ and $\beta_0 + \beta_1 \times mean(g)$ in equation (4). The three different sets of estimates in Table 4 correspond to different specifications of the aggregate disturbances: HP-filtered manufacturing inflation, commodity price inflation, and monetary policy shocks. Note first that on average firms are more willing to increase prices in response to adverse technology disturbances than lower prices in response to favorable shocks: the elasticity on positive shocks is close to -0.2, while that on negative shocks is close to -0.4 when the aggregate variables are at their steady-state means. I thus corroborate, although in a different environment, the results of Peltzman (2000) who finds that output prices are more likely to respond to cost increases, rather than decreases. I next compute the effect of a one standard deviation aggregate shock on the elasticities γ_t^N and γ_t^P (e.g., $\alpha_0 + \alpha_1 \times [mean(g) + std.dev.(g)]$). Notice that for all measures of aggregate disturbances, with the exception of manufacturing inflation rates, an increase in the size of the nominal disturbances reduces the elasticity of sectoral inflation to positive technology shocks by around 40% (e.g., from -0.23 to -0.14 for commodity price inflation), while increasing that on negative technology shocks by 30% (e.g., from -0.41 to -.59 for commodity price inflation). An increase in HP-filtered inflation rate itself increases the responsiveness to technology shocks in sectors with negative technology shocks, while leaving the elasticity in sectors with positive shocks unchanged.